




INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(51) International Patent Classification ⁷ : H04L 12/43, H04Q 11/00, H04L 12/56		A1	(11) International Publication Number: WO 00/69126
		(43) International Publication Date: 16 November 2000 (16.11.00)	
(21) International Application Number: PCT/GB00/01744		(81) Designated States: AE, AG, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BY, CA, CH, CN, CR, CU, CZ, DE, DK, DM, DZ, EE, ES, FI, GB, GD, GE, GH, GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MA, MD, MG, MK, MN, MW, MX, NO, NZ, PL, PT, RO, RU, SD, SE, SG, SI, SK, SL, TJ, TM, TR, TT, TZ, UA, UG, US, UZ, VN, YU, ZA, ZW, ARIPO patent (GH, GM, KE, LS, MW, SD, SL, SZ, TZ, UG, ZW), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, GW, ML, MR, NE, SN, TD, TG).	
(22) International Filing Date: 5 May 2000 (05.05.00)		<p>Published With international search report.</p>	
(30) Priority Data:			
99303672.2 11 May 1999 (11.05.99) EP 9919095.1 12 August 1999 (12.08.99) GB 9930373.7 22 December 1999 (22.12.99) GB			
(71) Applicant (for all designated States except US): BRITISH TELECOMMUNICATIONS PUBLIC LIMITED COMPANY [GB/GB]; 81 Newgate Street, London EC1A 7AJ (GB).			
(72) Inventors; and			
(75) Inventors/Applicants (for US only): COTTER, David [GB/GB]; 23 Moorfield Road, Woodbridge, Suffolk IP12 4JN (GB). LUCEK, Julian, Kazimierz [GB/GB]; 52 Wellesley Road, Ipswich, Suffolk IP4 1PL (GB).			
(74) Agent: WELLS, David; BT Group Legal Services, Intellectual Property Department, 8th Floor, Holborn Centre, 120 Holborn, London EC1N 2TE (GB).			

(54) Title: OPTICAL COMMUNICATIONS NETWORK



(57) Abstract

A node in an optical communications network receives optical packets at its input. It determines a priority assigned to each optical packet. When the optical packet is a transit packet destined for another node then, if the packet is determined to have a relatively higher priority, it is output from the node via a continuous-flow transmission path. Other transit packets having relatively lower priorities are sent via an alternative transmission path that may include queues or buffers resulting in a variable delay.

FOR THE PURPOSES OF INFORMATION ONLY

Codes used to identify States party to the PCT on the front pages of pamphlets publishing international applications under the PCT.

AL	Albania	ES	Spain	LS	Lesotho	SI	Slovenia
AM	Armenia	FI	Finland	LT	Lithuania	SK	Slovakia
AT	Austria	FR	France	LU	Luxembourg	SN	Senegal
AU	Australia	GA	Gabon	LV	Larvia	SZ	Swaziland
AZ	Azerbaijan	GB	United Kingdom	MC	Monaco	TD	Chad
BA	Bosnia and Herzegovina	GE	Georgia	MD	Republic of Moldova	TG	Togo
BB	Barbados	GH	Ghana	MG	Madagascar	TJ	Tajikistan
BE	Belgium	GN	Guinea	MK	The former Yugoslav Republic of Macedonia	TM	Turkmenistan
BF	Burkina Faso	GR	Greece	ML	Mali	TR	Turkey
BG	Bulgaria	HU	Hungary	MN	Mongolia	TT	Trinidad and Tobago
BJ	Benin	IE	Ireland	MR	Mauritania	UA	Ukraine
BR	Brazil	IL	Israel	MW	Malawi	UG	Uganda
BY	Belarus	IS	Iceland	MX	Mexico	US	United States of America
CA	Canada	IT	Italy	NE	Niger	UZ	Uzbekistan
CF	Central African Republic	JP	Japan	NL	Netherlands	VN	Viet Nam
CG	Congo	KE	Kenya	NO	Norway	YU	Yugoslavia
CH	Switzerland	KG	Kyrgyzstan	NZ	New Zealand	ZW	Zimbabwe
CI	Côte d'Ivoire	KP	Democratic People's Republic of Korea	PL	Poland		
CM	Cameroon	KR	Republic of Korea	PT	Portugal		
CN	China	KZ	Kazakstan	RO	Romania		
CU	Cuba	LC	Saint Lucia	RU	Russian Federation		
CZ	Czech Republic	LI	Liechtenstein	SD	Sudan		
DE	Germany	LK	Sri Lanka	SE	Sweden		
DK	Denmark	LR	Liberia	SG	Singapore		
EE	Estonia						

- 1 -

OPTICAL COMMUNICATIONS NETWORK

The present invention relates to a communications network, and in particular to a broadband optical network carrying optical packets.

5 Optical broadband networks have been deployed to meet the rapid increase in demand for bandwidth arising from the growth in the Internet and related data services. Conventionally, packets have been encapsulated for transmission using broadband network protocols such as ATM (Asynchronous Transfer Mode) or SDH (Synchronous Digital Hierarchy). However, it has been recognised that this
10 introduces a redundant layer of network processing and management and that it would be desirable to transmit packets directly over the optical transmission medium. For example, Cisco have proposed an "Optical Internetworking Roadmap" which describes five stages of evolution from data traffic being carried over SDH/SONET equipment through to full optical internetworking. The penultimate stage of this road
15 map has been reached with Cisco's introduction of architecture known as a "Dynamic Packet Transport" (DPT) technology. This uses a ring architecture in which packet routers are placed on two concentric counter-directional rings. The rings consist of optical fibre links, each of which terminate at the input and output ports of the routers. DPT operates by sending data packets in one direction and by sending
20 the corresponding control packets in the opposite direction on the other fibre.

According to a first aspect of the present invention, there is provided a method of operating a node in a communications network, comprising:

- (a) receiving an optical packet at the node,
- (b) determining a priority assigned to the optical packet,
- 25 (c) when the optical packet is determined to be a transit packet having a relatively higher priority, then outputting the packet from the node via a continuous flow transmission path,
- (d) for at least some packets determined to have a relatively lower priority, transmitting the said packets through the node via an
30 alternative transmission path including one or more queues.

The present invention provides a method of operating a node, which allows express packets, that is high priority transit packets, to flow past any intermediate routers between the source and destination. This gives the advantages of reduced

- 2 -

traffic load at the router, greater throughput, reduced delay, and reduced delay variation. At the same time, the nodes remain capable of implementing queuing and statistical multiplexing of packets where required. Using the invention, express multicast control packets may be transmitted to implement time-critical network management operations, such as support recovery. The fixed speed-of-light latency of express packet transfer also benefits time-critical applications such as distributed database concurrency. It also enables advanced latency-management techniques.

A continuous-flow transmission path is one which transmits packets in a steady, continuous fashion without using queues or other variable delays. The continuous-flow transmission path may be entirely in the optical domain. Alternatively, part or all of the continuous-flow transmission path may be in the electrical domain. In this case, electrical-optical conversion will be carried out at the input and output to the node.

According to a second aspect of the present invention, there is provided a node for connection in a communications network, the node comprising a continuous-flow transmission path connecting the input to the node to the output from the node, an alternative path connecting the input to the node to the output from the node and including one or more packet queues, and

control means arranged to determine the priority assigned to a transit packet to direct a transit packet having a relatively higher priority via the continuous-flow transmission path, and to direct a transit packet having a relatively lower priority via the alternative transmission path.

According to a third aspect of the present invention, there is provided a method of operating a node in a communications network, comprising

- (a) receiving a packet in the optical domain at an input to the node,
- (b) creating a copy of the packet,
- (c) directing at least some of the transit packets received at the input of the node via a continuous-flow transmission path, to the output of the node.

Preferably, the node includes a switch connected to its output, and the method includes selectively ejecting at least some of the packets from the continuous-flow transmission path. This aspect of the invention provides a new node architecture termed by the inventors the "copy-eject-and-insert" (CEI) architecture.

- 3 -

This architecture facilitates the transmission of high-speed multicast packets, and is also particularly suitable for use with the packet prioritisation scheme of the preceding aspects of the invention. In a conventional "drop-and-insert" (D&I) architecture, packets are extracted from the transmission path for the purpose of receiving them at a node. In the case of a CEI node embodying the invention, by contrast, packets are copied but not extracted from the transmission path when they are to be received by the node. All of the incoming traffic may be copied to the node, for example by means of a passive optical coupler. If an incoming packet has reached its destination or is otherwise intended to terminate at the node, then the packet may also be ejected from the transmission path, for example by means of a crossbar switch at the output of the node. When a multicast packet is received at the node, it may both be read from the copied packet and concurrently transmitted onwards via the continuous transmission path.

Systems embodying the present invention will now be described in further details, by way of example only, with reference to the accompanying drawings in which:

Figure 1 is a schematic showing a prior art network;

Figure 2 is a schematic showing a first example of a network embodying the present invention;

Figure 3 shows the data flows through a node in the network of Figure 2;

Figure 4 is a schematic showing the architecture of a prior art node;

Figure 5 is a schematic showing the architecture of a node embodying the present invention;

Figure 6 is a schematic showing an alternative node architecture;

Figures 7i to 7iv show alternative configurations for the node of Figure 5;

Figure 8 illustrates the substitution of a high priority packet or a low priority packet;

Figure 9 shows data files within a node for different packet types;

Figure 10 shows the format of an optical packet;

Figure 11 shows the use of dual addresses;

Figure 12 shows a ring-wrap operation;

Figure 13 shows contention at a node output;

Figure 14 shows logic and synchronisation circuits;

- 4 -

Figure 15a shows the structure of a node used to switch wavelength division multiplexed optical packets;

Figure 15b shows the format of a packet handled by the node of 15a; and

Figures 16a to 16e show crossbar switch structures.

5 Figure 1 shows a prior art architecture developed by Cisco and known as the DPT (Dynamic Packet Transport) ring architecture. In this scheme, Internet routers are placed on two concentric counter-directional rings. The rings consist of a sequence of optical fibre links, each of which terminates at the input and output ports of the routers. One of the rings is known as the inner ring and the other as the
10 outer ring. Data packets are sent in one direction and corresponding control packets are sent in the opposite direction on the other fibre. This IP (Internet Protocol) ring network is designed to provide a number of features. In particular, it supports statistical multiplexing of packets with no provision of point-to-point connections or use of dedicated bandwidth for interconnection of routers or for protection. It is also
15 designed to support packet prioritisation and to offer multiple levels of queuing and scheduling and to support both unicast and multicast transmissions. In the network of Figure 1, by contrast with the network described below with reference to Figure 2, the signal transmission path is interrupted at each node.

 Figure 2 shows a network embodying the present invention. A
20 communications network 1 comprises concentric outer and inner rings 2.1, 2.2. The rings carry optical packets. A number of nodes 3a to 3g are connected to the rings. In this example, each of the nodes 3a to 3g is an Internet Protocol router. In operation, one of the routers, for example 3a, outputs an optical packet addressed to another of the routers, for example 3d, onto one of the dual concentric rings. The
25 optical packet output by node 3a carries the network address of node 3d in an address field in the packet header. The packet passes around the ring. At each intermediate node, the packet address is compared with the address of the respective node. If the packet is not addressed to a given intermediate node, then the packet passes on to the next node, and so on. At the destination node, the address is read
30 and the packet is recognised as being intended for receipt at that node. The packet, in the case of a unicast transmission is removed from the network at that node. It may then, for example, be converted to the electrical domain for onwards

- 5 -

transmission to a customer terminal on an electronic network connected to the respective node.

Figure 3 shows the flow of received and transmitted data packets at one of the nodes on one of the fibre rings. An incoming data packet is (i) sent to host receive queue (if it has reached its destination) or (ii) is passed over for onward transmission. A multicast packet may both be sent to the host receive queue and also be forwarded for onward transmission. Packets for onward transmission are treated according to their priority: a low-priority packet is passed to a transit queue. An express high-priority, packet is passed immediately to the outgoing path without queuing. When capacity on the ring allows, a packet from one of the transmit queues may be passed to the outgoing path. Multiple transmit queues may be used to manage prioritisation of outgoing traffic from the node.

In this example, the optical fibre rings and the nodes are configured to provide a continuous optical path for express packets.

In a continuous transmission path signals stream into and out from the path at a continuous and uniform rate. The path may be a continuous optical transmission path. The continuous optical transmission path may contain a substantially fixed delay as in Figure 3. In the case of a continuous transmission path in the electrical domain the path may contain storage elements such as an FIFO (first-in first-out) buffer. This differs from conventional systems, however, in this case, in that in normal operation, the output from the buffer is continuous and regular, so that the input and output of the buffer occur at substantially the same rate. In conventional prior systems, a buffer is used to hold transit packets in a queue which is serviced at rates and/or times which are dependent on factors such as packet priorities, queue length, status of contending queues etc.

A further function carried out by the control circuits is to determine whether the time-to-live (TTL) of a packet has expired. If the time-to-live is determined to have expired, then again a control signal is sent to the switch to cause such a packet to be ejected from the onwards transmission path. When one of the nodes generates and outputs an optical packet onto the ring, a value is written in a phase field "PH" in the packet header. One of the nodes functions as a master node. The master node outputs multicast express control packets that travel around each ring to inform each node of the current phase value. The phase value is regularly updated, at a rate of at

- 6 -

least once per round-trip time for speed-of-light travel around the ring. The updating of the phase value is carried out by the master node incrementing the value. For example, in the case of one control packet per round trip of the ring, the phase value is incremented by 1 (modulo 3). Each ring periodically receives and stores the
5 current phase value as indicated by one of the multicast express control packets from the master node. When any other express packet is received at the node the control logic in the node determines whether the value in the PH field of the received express packet is such that:

$$(\text{current_phase} - \text{PH}) \bmod 3 > 1,$$

- 10 If this inequality is satisfied, then the packet is expired and should be purged from the ring. This may be the case, for example, because the packet is addressed to a node which is currently malfunctioning or because there is an error in the address of the packet.

Figure 4 shows the traditional 'drop-and-insert' (D&I) architecture that has
15 been widely considered previously for optical packet networks (e.g. J R Sauer, M N Islam and S P Djaili, "A soliton ring network", Journal of Lightwave Technology, vol. 11, no. 12, December 1993, pp. 2182-2190). The node contains a 2x2 optical crossbar switch. When the switch is in the bar (straight-through) position: (i) an optical packet on the incoming fibre is 'dropped' (i.e. switched to the node receiver);
20 or (ii) an optical packet created by the host transmitter is 'inserted' (i.e. switched to the outgoing fibre); or both (i) and (ii) occur simultaneously. A limitation of the D&I node architecture is that when a packet is received by the node it is physically removed from the optical path between the incoming and outgoing fibres. The only method for multicasting a packet is, at each node, to drop the packet and then
25 subsequently reinsert it when sufficient vacant capacity on the ring becomes available. Therefore this architecture cannot support multicasting of express (high-priority) transit packets.

Figure 5 shows a preferred form of the new node architecture that we have invented. Instead of 'drop-and-insert' (D&I), this new architecture provides the
30 function of 'copy-eject-and-insert' (CEI). In the case of a D&I node, packets are extracted from the transmission path for the purpose of receiving them at a node. In the case of a CEI node, packets are copied but not extracted from the transmission path for the purpose of receiving them at a node. All incoming traffic is copied to the

- 7 -

node, for example by means of a passive optical coupler as shown in Figure 5. In use, an incoming fibre from one of the optical rings passes to a coupler 5.1. From the coupler 5.1, one optical path passes through a further length of fibre 5.3 providing a fixed delay to an optical crossbar switch 5.4. An outgoing fibre
5 connected in the fibre ring passes from one of the output ports of the crossbar switch. The other output from the coupler 5.1 is split again, for example, using a further coupler 5.2. One branch of the output from this further coupler is used to provide a copy of the optical packet on the incoming fibre. This copy may be passed, for example, to the host receive queue. The other branch of the output of
10 the further coupler passes to control circuits. These control circuits may, for example, read a header carried with the optical packet, and carry out associated logic operations. A control output is generated by the control circuits and passes to the crossbar switch. If the control circuits determine that the packet is addressed to the respective node, then the crossbar switch is set to the cross state in order to eject
15 the packet (in the case of a unicast transmission). If an incoming packet has reached its destination or is otherwise intended to terminate at the node, the packet is also ejected from the ring, for example by means of the crossbar switch as shown in Figure 5. Simultaneously, or whenever sufficient vacant capacity on the ring become available, a packet may be inserted onto the ring by the node. The control logic and
20 synchronisation sub-systems use information contained in the packet (for example in the packet header) and other information (for example the status of transmit queues and the status of control flags) to operate the crossbar switch. Figure 5 shows a fixed optical delay, which allows sufficient time for the operation of these control logic, synchronisation sub-systems and switch before the arrival of the packet at the
25 switch. As in the D&I architecture, an express (high priority) transit packet is passed directly to the outgoing fibre. However, unlike D&I, the CEI architecture allows express transit packets to be multicast, because they are simultaneously copied without delaying their onward passage.

Figure 14 shows the control logic subsystem in more detail. It includes
30 optical logic stages 141 and combinatorial electronic logic gates 142. The four optical blocks (synch, address recognition PH reader and Unicast/Multicast) each have copies of the optical packet to the input. The output from the synchronisation block is an optical path to each of the other three optical blocks. The synch block

- 8 -

may be based on one of the self-synchronisation techniques we have described in our patent EP-B-687370 (eg. using pulses separated by 1.5 bit periods input to an optical AND gate). The address recognition block may be based on the technique we have described before in that patent. Addresses are coded using specially selected binary words, and recognised by inputting address and target word to an optical AND gate. The output from the AND gate is converted to give an output from the block that is an electrical binary signal that says packet ADDRESSEE matches/does not match local address. The PH may be two optical AND gates, each having the optical packet as one input and a synch pulse as the other input - this synch pulse timed to overlap with one of the PH bits. The output from PH reader block is two parallel electrical binary signals - each denoting one of the PH bits. The UNI/MULTI reader may be one optical AND gate, having the optical packet as one input and a synch pulse as the other input - this synch pulse timed to overlap with the UM bit. The output from UNI/MULTI reader block is one electrical binary signal - denoting unicast/multicast. These electrical signals, together with signals from the queue status then pass to the high-speed electronic logic part. This part performs the logic which is set out in the table below describing the action of the node for various types of packet (e.g. multicast low-priority transit packet, etc). The output from this electronic logic is a binary electrical signal to set the 2x2 cross bar optical switch configuration.

Each stage of this logic has to operate within a time shorter than the shortest packet. e.g. for a 50 byte packet at 100 Gbit/s = 4 ns. Since the optical stage and electronic logic stage are arranged in pipeline, each stage would have to take less than the minimum time (eg 4 ns). The optical stage is ultrafast - e.g. using four-wave mixing in semiconductor optical amplifier, or TOAD device, etc. to implement the optical AND gate. Since the electronic logic required is simple it is possible to construct suitable fast circuit using hard-wired combinatorial logic.

Figure 16a shows the design of a simple crossbar switch. This includes a region 161 of lithium niobate. When an electrical control signal is applied to a gate on the lithium niobate, then the refractive index of the lithium niobate changes, so that the optical signal is coupled across from one optical path to the other, the "cross" state of the switch.

Figure 16b shows an alternative design for a crossbar switch. In this case, the switch uses a Mach Zehnder interferometer configuration with a control element

- 9 -

162 in one branch of the interferometer. A phase change is generated by the control element 162 when the control signal is applied, to switch the output from one optical output port to the other. Both of these designs suffer significant cross talk: that is, when an optical signal is switched to one of the ports, a significant level of optical power, for example 20 or 30 percent of the input power, is present at the other output port.

To overcome this problem, each crossbar switch in the nodes described above may be replaced by a network of four such switches in the configuration shown in Figure 16c. The switches are cross-connected as shown. The control signal c is applied to the first pair of switches in the signal path, and the logical complement of the control signal, \bar{c} is applied to each of the second pair of switches in the signal path. Logically, the switching effect of such a network is equivalent to that of each the single switches of Figure 16a or 16b. However, the crosstalk is reduced by an order of magnitude. Since the crossbar switch required only two inputs, one for the insert path and one for the straight through path, and only two outputs, one for the eject path and one for the output path, it can be shown that the full network shown in figure c is not required. Instead, the full crossbar switching function with the benefit of crosstalk reduction can be achieved using a pair of cross bar switches connected in series as shown in Figure 16d.

Figure 16e shows a further alternative design also offering improved cross talk performance. This uses a pair of passive optical couplers 163, 164 with an electro-absorption modulator (EAM) 165 connected between them. The copy optical signal is taken from the first of the optical couplers 163 and the insert optical signal is added at the second of the optical couplers 164. In the absence of a control signal c , the EAM 165 is in its transparent state and the optical packet passes straight through from the input path to the output path. When a packet is to be ejected then the control signal c switches on the EAM 165. The ejected packet may then be replaced by further packet added on the INSERT optical path of the coupler 164.

Figure 6 shows an example of the CEI architecture in an alternative form. Here there are two crossbar switches, SW1 and SW2. (For clarity, the control logic and synchronisation sub-systems are not shown.) In this case not all incoming packets are copied; instead only those that arrive whilst SW1 is in the cross state will be copied. In this case an express multicast packet may be copied and

- 10 -

immediately reinserted onto the ring if switches SW1 and SW2 are both placed in the cross state. However in this case the express multicast packet is delayed by the transit time of the feedback loop from SW1 to SW2 and back again to SW1. The alternative form of the CEI architecture shown in Figure 6 has certain drawbacks:

- 5 First, the forced additional delay for express multicast packets increases the probability of contention at SW1 (described further below) unless special preventative measures are used (such as, for example, providing a sufficient time guard band following an express multicast packet, with the resulting penalty of reduced network throughput and higher complexity). Second, the architecture shown in Figure 6
10 requires two optical switches, rather than one.

In the case of the D&I architecture (Figure 4), the extraction of packets from the transmission path for the purpose of receiving them at a node is necessarily a selective process in the optical domain, and may be performed for example by an optical crossbar switch. In the preferred form of the CEI architecture (Figure 5), the
15 copying of packets from the transmission path for the purpose of receiving them at a node is non-selective in the optical domain, and may be performed for example by a passive optical coupler.

In the CEI architecture the continuous transmission path for express transit packets is preferably a continuous optical transmission path, but not necessarily so,
20 as shown in Figures 7 (i)-(iv). Figure 7(i) shows a preferred arrangement with a continuous optical transmission path for express transit packets. Figure 7(ii) shows a transponder (a regenerator using optical-electrical and electrical-optical conversion stages). Figure 7(iii) shows a case where the copy function is performed in the electrical domain. Figure 7(iv) shows a case where the copy, eject and insert
25 functions are all performed in the electrical domain. In this case the switch can be a 1x2 electrical switch, and as in the case of optical switching, the channel left open circuit is in effect 'ejected'.

Figure 7(iv) is distinguished from conventional nodes by a 'continuous transmission path'. In the continuous transmission path signals stream into and out
30 from the path at a continuous and uniform rate.

An important feature of an optical packet network with different priority classes is the ability to reallocate the time occupied on the ring by a transit packet to another packet with higher priority. The reuse of certain time slots using a D&I

- 11 -

optical node architecture is previously known. For example, the previously cited paper of Sauer et al describes a D&I optical node in which a slot made vacant by dropping a packet (i.e. removing from the optical ring a packet when the packet destination address matches the node address) can be immediately reused by inserting a transmit packet. Here, distinctively, we reallocate any time slot by ejecting an incoming packet from the ring so as to vacate the time slot, and immediately or simultaneously insert into the same time slot another packet taken from one of the transit or transmit queues. Using the CEI architecture, the steps required to reallocate to a packet of higher priority a time slot currently occupied by an incoming transit packet are: copy the incoming transit packet; determine its priority; if appropriate for reallocation, then do {eject the packet from the ring; deliver the packet copy to the transit buffer; transmit another packet of higher priority reusing the same time slot}. This is illustrated in Figure 8. Notice that the CEI architecture is not essential for this process. For example, using a D&I architecture the steps required would be: determine the priority of an incoming transit packet; if appropriate for reallocation, then do {drop the packet; store the packet for later transmission; and transmit another packet of higher priority reusing the same time slot}. Although reference is made here to 'slot', this feature is equally applicable to unslotted systems in which packets arrive asynchronously, and may be of variable length.

Figure 9 shows the flow of incoming and outgoing packets in the optical CEI node of Figure 5, so as to realise the architecture of Figure 3. All packets are assumed to be in one of two priority classes: low and high (express). Packets may also be unicast or multicast, and they may be data packets (with a payload of user data) or control packets (uniquely for network control purposes). All incoming packets are copied, whereupon they may be allocated to various input queues or else discarded. The operations at the node to handle different types of legitimate (e.g. non-expired) incoming packet are as follows:

- Unicast
- Express transit packet: The crossbar switch is put in the bar (straight through) state so that the packet is passed directly to the outgoing fibre. The packet copy is discarded.

- 12 -

- Express packet addressed to node: The crossbar switch is put in the cross state to eject the packet from the ring. The packet copy is delivered to the input queue, and subsequently sorted according to whether it is a data or control packet.
- 5 • Low-priority transit packet: The crossbar switch is put in the cross state to eject the packet from the ring. The packet copy is delivered to the low-priority transit buffer.
- Low-priority packet addressed to node: The crossbar switch is put in the cross state to eject the packet from the ring. The packet copy is delivered
10 to the input queue, and is subsequently sorted according to whether it is a data or control packet.
- Multicast
 - Express packet (different source): The crossbar switch is put in the bar state so that the packet is passed directly to the outgoing fibre. The
15 packet copy is delivered to the input queue, and is subsequently sorted according to whether it is a data or control packet intended for delivery to the node, otherwise it is discarded.
 - Express packet (node is the source): The crossbar switch is put in the cross state to eject the packet from the ring. The packet copy is
20 discarded.
 - Low-priority packet (different source): The crossbar switch is put in the cross state to eject the packet from the ring. The packet copy is delivered
25 to low-priority transit buffer. The packet copy is also delivered to the input queue, where it is subsequently sorted according to whether it is a data or control packet intended for delivery to the node, otherwise it is discarded.
 - Low-priority packet (node is the source): The crossbar switch is put in the cross state to eject the packet from the ring. The packet copy is
30 discarded.
- 30 Packets for transmission from the node are selected from the low-priority transit buffer or the transmit queues according to the current status of queue depths, appropriate rate controls, fairness algorithm, etc. Subject to these controls, packets are transmitted as vacant capacity becomes available on the ring (either vacant

- 13 -

capacity on the incoming fibre or new vacant capacity created by the ejection of packets from the ring by the node itself).

Each packet carries an appropriate MAC protocol header, designed according to the principles set out earlier (Figure 10 shows a suggested example). This MAC header can consist of two parts: one part (which we call the 'optical MAC header') contains the minimum information needed for the processing functions that must be carried out at high speed 'on the fly' to enable the operations listed above. The second part of the MAC header contains information needed for other Layer 2 functions.

10 In the example shown in Figure 10, SS (self-synchronisation) is a 2-bit field used for timing recovery. ADDRESSEE is a 10-bit field used to indicate which node should strip the packet from the ring. In the 'Spatial Reuse Protocol' (SRP) described by Cisco (part of the Dynamic Packet Transport), packets are stripped from the ring by the destination node in the case of unicast packets. It is proposed here that for 15 the purpose of routing in the high-speed network, the optical MAC header will not include both the source and destination addresses, but instead will contain a single address field ADDRESSEE. For the transmission of a unicast packet, ADDRESSEE will be set to the destination address. For the transmission of a multicast packet, ADDRESSEE will be set to the source address. A node is required to eject a packet 20 from the ring if ADDRESSEE matches the address of the node. UM (unicast/multicast) is a 1-bit field used to indicate whether the packet is unicast or multicast. PH (phase) is a 2-bit field with a dual purpose; as described more fully below, PH indicates the priority of the packet and also contains a simple time-to-live mechanism. P (parity) is a 1-bit field used to set the parity of the optical MAC header 25 (approximately 2 bytes), and will be used in combination with 'on the fly' parity checking to provide some header integrity.

The reading of the optical MAC header must be performed 'on the fly' using very simple high-speed logic operations. For ultra-high-speed optical implementation, processing based on simple optical logic gates, such as AND, will be used. Various 30 methods of timing recovery based on self-synchronisation are described in our patent EP-B-687370. For example, the field SS could consist of 2 return-to-zero format optical pulses separated by 1.5 bit periods, and self-synchronisation would be based on the output from an optical AND gate whose first input is a copy of the packet and

- 14 -

the second input is a further copy of the packet delayed by 1.5 bit periods with respect to the first. The output from the AND gate is an optical pulse in precise synchronism with the start of the packet. This pulse (which we term here the 'timing pulse') can then be used for a variety of operations on the fly. The timing pulse may
5 be used in the process of determining whether there is a match between ADDRESSEE and the address of the node, as described in our patent EP-B-687370. The timing pulse may also be used in combination with optical AND gates to read the fields UM and PH.

The dual ring structure is resilient because, in the event of an isolated fibre
10 break or node failure, the nodes can perform the 'ring wrap' operation illustrated in Figure 12 [as described for example in Cisco's recently published white paper on Dynamic Packet Transport Technology and Applications]. For proper operation of the control mechanisms which allow the network to reconfigure and re-establish after ring wrap, it is clearly necessary for the nodes to distinguish between those incoming
15 packets which are travelling on their 'correct' ring and those that have been wrapped onto their 'wrong' ring. Here we propose the technique of 'dual addressing' to avoid the need for an additional header field to indicate the 'correct' ring for each packet. Dual addressing is illustrated in Figure 11. Rather than provide each node with a single address, it is proposed to provide each node with two addresses, one for each
20 of the rings. In this case, the transmitter of a packet uses the appropriate value of ADDRESSEE corresponding to the ring that is being used. In the event of ring wrap, the usual rule for packet stripping is simply followed (a node is required to eject a packet from the ring if ADDRESSEE matches the address of the node), without the need to read further header fields to check whether the packet has been wrapped
25 onto the 'wrong' ring.

Some network control operations require point-to-point signalling between adjacent nodes. This can be done in a number of ways: creation of an independent control packet; 'piggy-back' technique such as overwriting certain fields in a newly created data packet or in a low-priority transit packet; or out-of-band signalling. Out-
30 of-band signalling could be performed in time guard bands between packets.

A potential limitation of the CEI node architecture is the possibility of contention at the eject-and-insert switch. This is illustrated in Figure 13. Contention may arise when an incoming packet B arrives when the switch is in the cross position

- 15 -

and the node is in the process of inserting a packet A. In the absence of an optical buffering mechanism, various possibilities for contention resolution are available including switch over to the bar state to allow B to pass, thus forcing the ejection of part of A. Resend A later. The recipients of A will recognise that the packet has
5 been truncated, and discard it.

A node embodying the present invention may also be used in a network employing wavelength division multiplexing. Figure 15a shows one example of such a node. A wavelength division multiplexed packet on the incoming fibre passes to a wavelength division demultiplexer 15. Suitable wavelength division multiplexers are
10 available commercially and may be formed, for example, from wavelength-selective optical couplers or from a fibre Bragg grating. The demultiplexer 151 outputs different wavelength channels on different output fibres. For ease of illustration, only four wavelength channels, $\lambda_1 - \lambda_4$ are shown, but in practice, a greater number of wavelength channels, for example 8, 16 or 32 channels may be used. The optical
15 signal on each wavelength channel is passed through an opto-electronic converter, such as a photodiode, and is converted to a signal in the electrical domain. The outputs from the opto-electronic converters are tapped to provide the copy of the packet. The copy may be passed to the host receive queue and to control/synchronisation circuits, as described previously in relation to the other
20 embodiments. The outputs of the opto-electronic converters are also connected via an electrical continuous-flow path to the input of an electronic switch 152. A parallel datastream of signals to be inserted into the flow path are also input to the switch 152. The parallel outputs of the switch 152 pass through electro-optic converters and are converted back into the optical domain at different respective wavelengths
25 $\lambda_1 - \lambda_4$. The signals at the different wavelengths are combined onto the outgoing fibre by an optical multiplexer 153. Figure 15b shows the format of the optical packets as received on the incoming fibre and as output on the outgoing fibre. In this example, 32 wavelengths are used and the packets are 32bits wide in the wavelength domain. Bit positions 1, 33, 65,... are received at wavelength λ_1 , bit
30 positions 2, 34, 66,... at wavelength λ_2 , and so on. In the node, in the input fibre prior to demultiplexing of the WDM packet, the optical signals may be passed through a dispersion-compensating device, such as a length of dispersion-

- 16 -

compensating fibre, to realign the different wavelength channels in the time domain so as to correct for the effects of group-velocity dispersion in the optical network.

- 17 -

CLAIMS

1. A method of operating nodes in a communications network comprising:
 - (a) receiving an optical packet at the node;
 - 5 (b) determining a priority assigned to the optical packet;
 - (c) when the optical packet is determined to be a transit packet having a relatively higher priority, then outputting the packet from the node via a continuous flow transmission path;
 - (d) for at least some packets determined to have a relatively lower priority,
 - 10 transmitting the said packets through the node via an alternative transmission path including one or more queues or buffers.
2. A method according to Claim 1, in which all of the continuous flow transmission path is in the optical domain.
- 15 3. A method according to Claim 1 or 2, in which the priority of a packet is determined from a flag carried with the packet.
4. A method according to any one of the preceding claims, including creating a
20 copy of the packet, when the packet has a relatively higher priority directing one of the original of the packet and the copy of the packet via the continuous-flow transmission path, and when the packet has a relatively higher priority directing the other of the original of the packet and the copy of the packet via the alternative transmission path;
- 25 5. A node for connection in a communications network, the node comprising a continuous-flow transmission path connecting the input of the node to the output of the node,
an alternative transmission path connecting the input of the node to the
30 output of the node and including one or more packet queues or buffers; and control means arranged to determine the priority assigned to a transit packet, and to direct a transit packet having a relatively higher priority via the

- 18 -

continuous-flow transmission path and to direct a transit packet having a relatively lower priority via the alternative transmission path.

6. An optical communications network having a ring topology and including one
5 or more nodes according to claim 5.

7. A method of operating a node in a communications network, comprising
a) receiving a packet in the optical domain at an input to the node,
b) creating a copy of the packet,
10 c) directing at least some of the transit packets received at the input of
the node via a continuous-flow transmission path, to the output of the node.

8. A method according to claim 7, in which the node includes a switch connected to
its output, and the method includes selectively ejecting at least some of the packets
15 from the continuous-flow transmission path.

9. A method according to any one of claims 1 to 4 and 7 to 8, in which the optical
packet received in step (a) is a wavelength division multiplexed packet, the method
including demultiplexing the optical packet in the wavelength domain.

20
10. A node according to claim 5, further comprising a wavelength demultiplexer
connected to the input to the node and arranged to demultiplex a wavelength-
multiplexed packet received at the input to the node, and a wavelength multiplexer
connected to the output of the node and arranged to multiplex a plurality of
25 wavelength channels to form a wavelength division multiplexed optical packet for
transmission from the output of the node.

11. An optical switching system for use in a node according to claim 5 or 10, the
optical switching system comprising first and second optical crossbar switches
30 connected in series, means for applying a control signal to the control input of the
first optical crossbar switch, and means for applying a complementary control signal
to the control input of the second optical crossbar switch.

- 19 -

12. An optical switching system for use in a node according to claim 5 or 10, the optical switching system comprising a first optical coupler connected in series with a second optical coupler, and an electro-optical absorber connected between an optical output of the first optical coupler and an optical input of the second optical coupler.




Figure 1




Figure 2

3/19




Figure 3




Figure 4

5/19




Figure 5

6/19




Figure 6

7/19

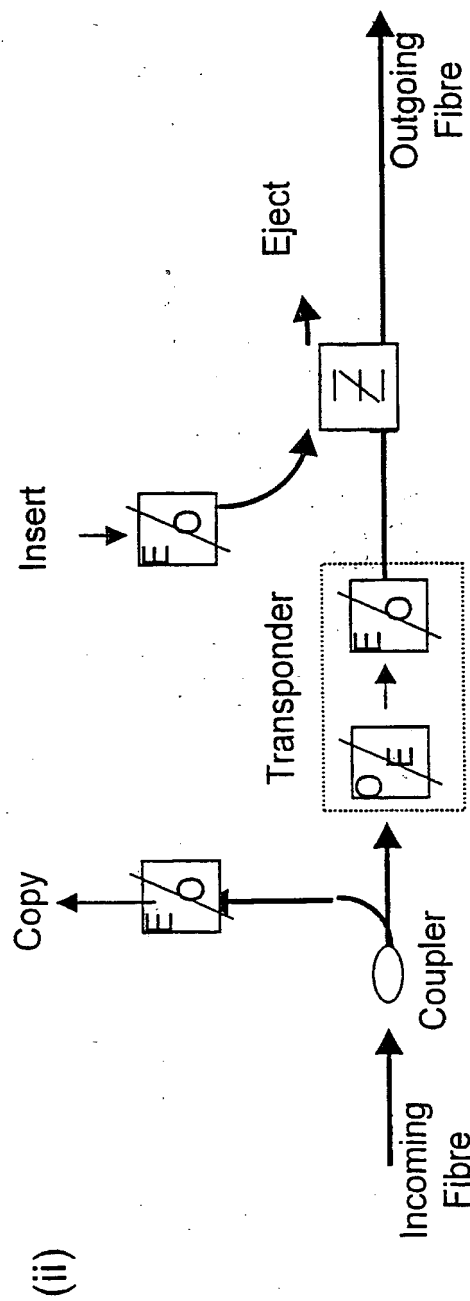

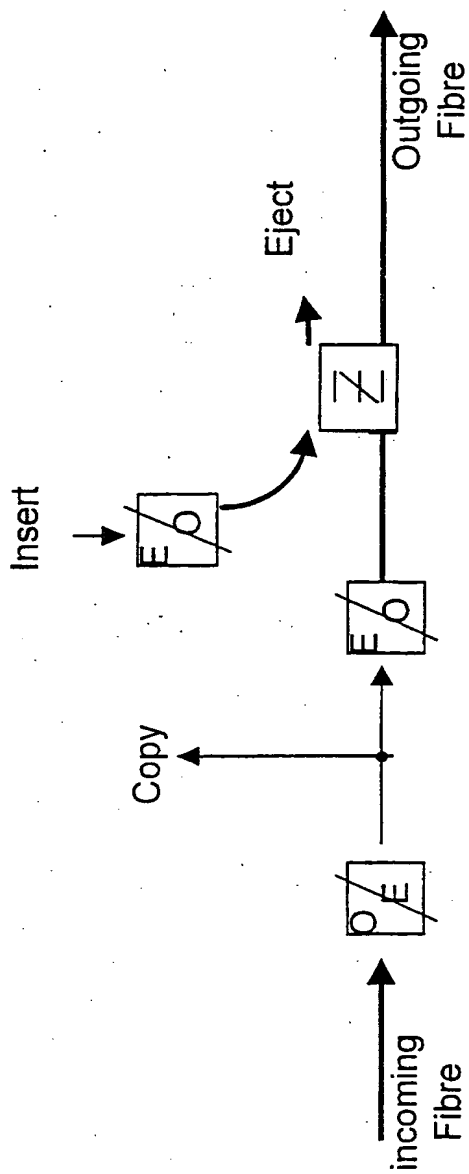


Figure 7(i), 7(ii)

(iii)



(iv)





Figure 7(iii), 7(iv)

9/19



B = received packet or lo-priority transit packet

D = packet from transmit queue or lo-priority transit buffer

Figure 8

10/19

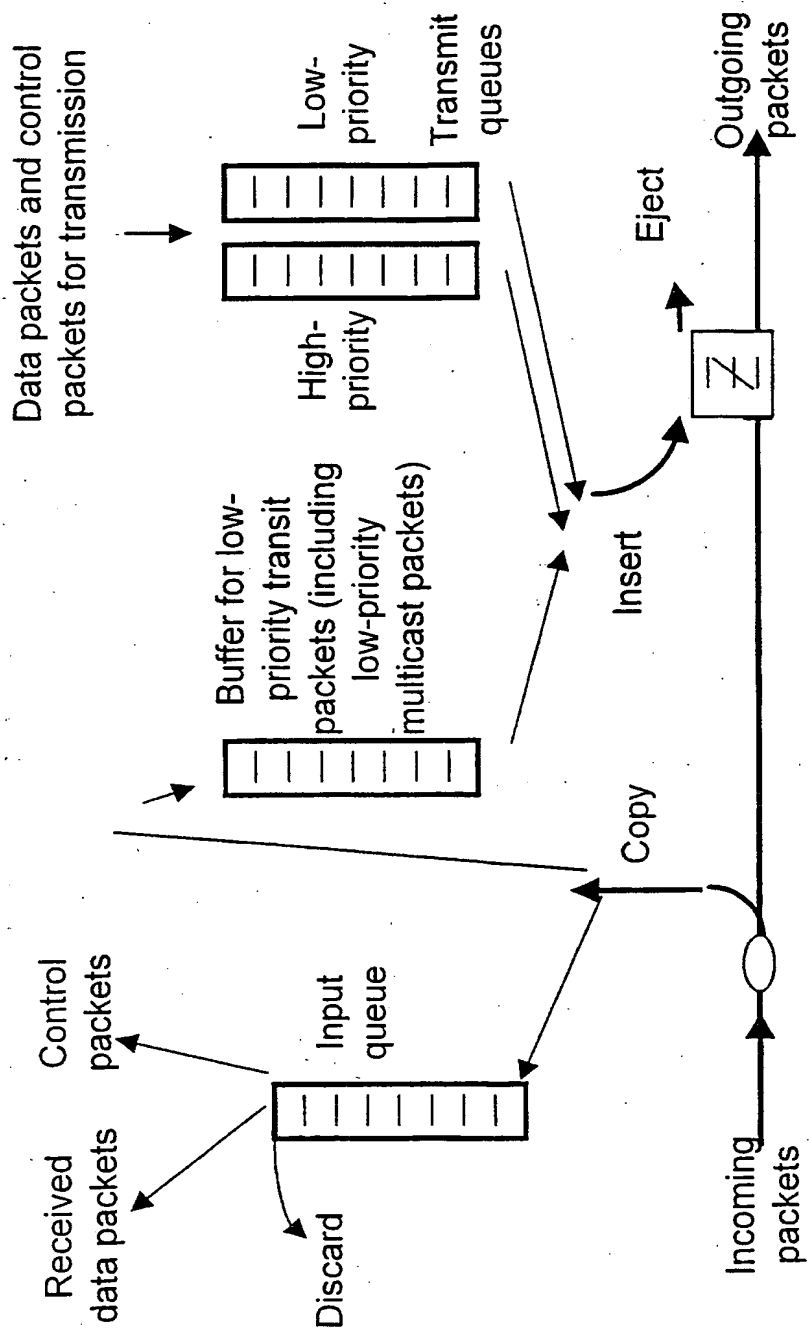


Figure 9

11/19




Figure 10




Figure 11




Figure 12

14/19




Figure 13

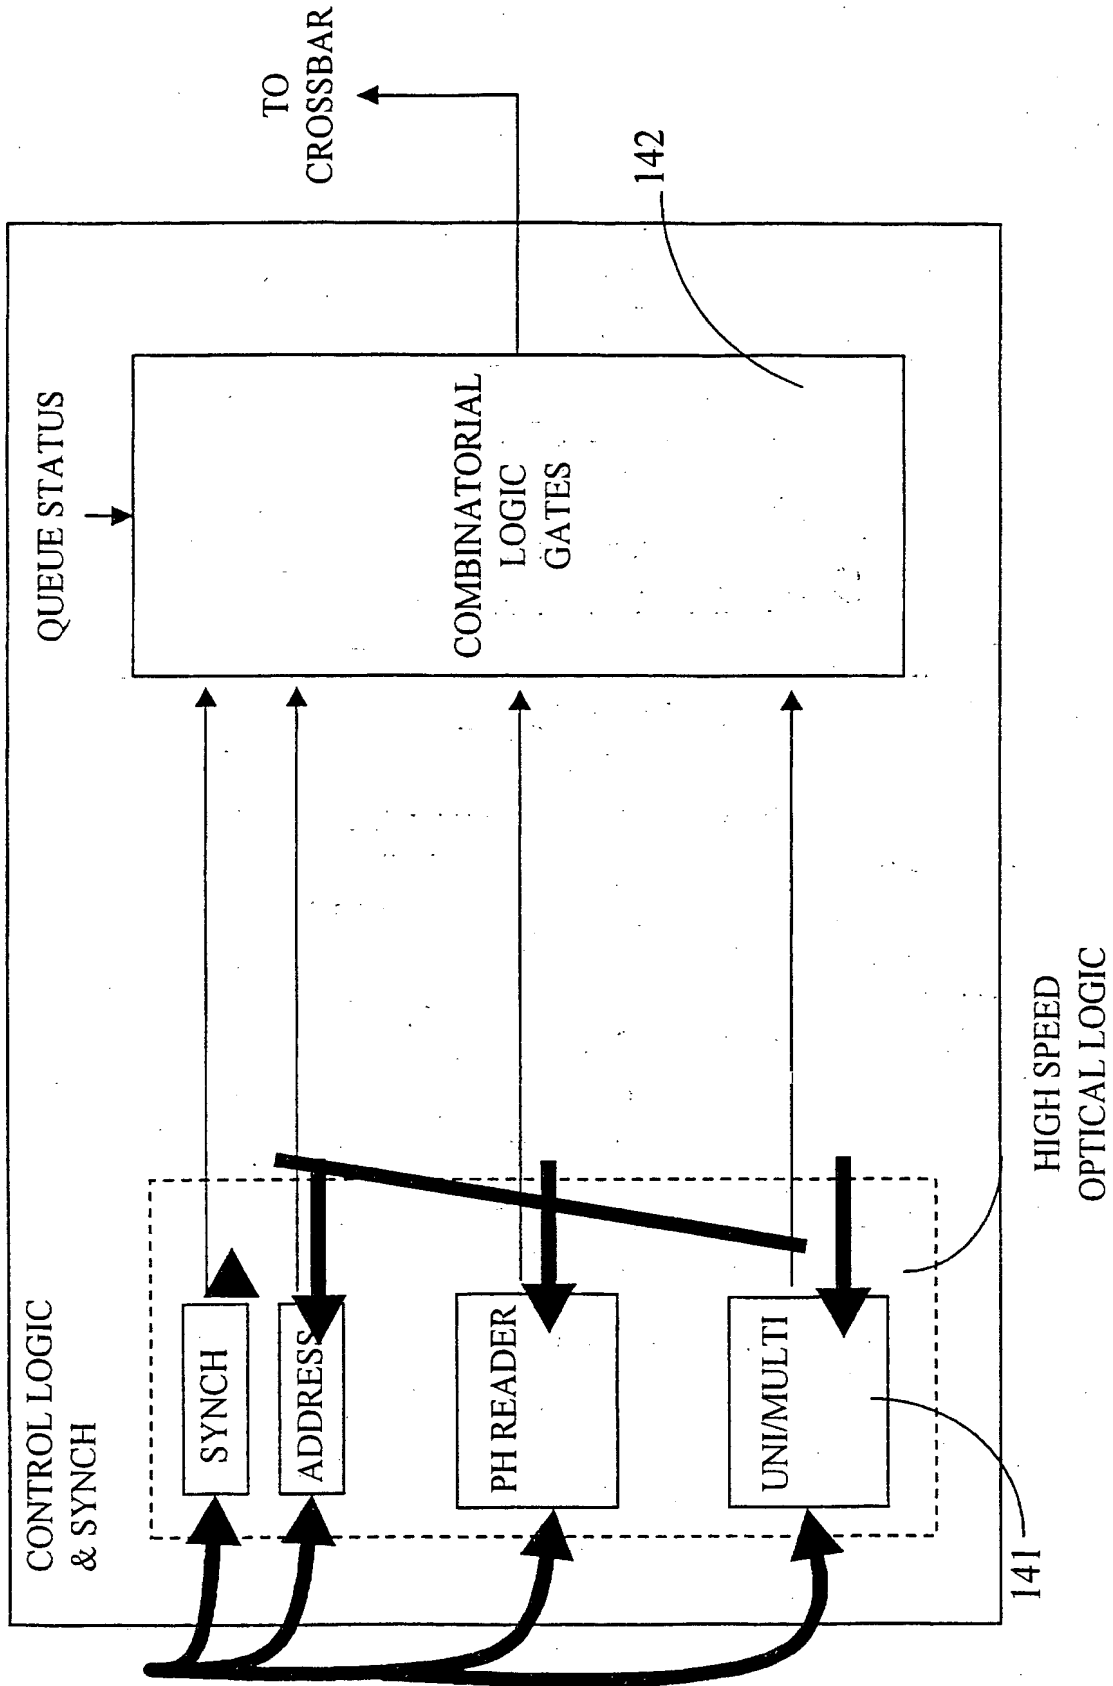


Figure 14

16/19




Figure 15a







Figure 15b

18/19

Figure 16



19/19



INTERNATIONAL SEARCH REPORT

International Application No

PCT/GB 00/01744

A. CLASSIFICATION OF SUBJECT MATTER

IPC 7 H04L12/43 H04Q11/00 H04L12/56

According to International Patent Classification (IPC) of the both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

IPC 7 H04L H04Q

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practical, search terms used)

EPO-Internal, INSPEC

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category *	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	EP 0 462 349 A (IBM) 27 December 1991 (1991-12-27)	1-8
A	page 2, line 1 - line 5 page 2, line 47 - line 55 page 4, line 46 - line 48 page 5, line 35 - page 6, line 37 page 7, line 2 - line 56 figure 4	9-12
X	SAUER J R ET AL: "A SOLITON RING NETWORK" JOURNAL OF LIGHTWAVE TECHNOLOGY, vol. 11, no. 12, 1 December 1993 (1993-12-01), pages 2182-2190, XP000422682 ISSN: 0733-8724 cited in the application	11
A	page 2183, paragraph II - page 2184; figures 1-3	1-10, 12
	--- -/-	

☒ Further documents are listed in the continuation of box C.

☒ Patent family members are listed in annex.

* Special categories of cited documents:

"A" document defining the general state of the art which is not considered to be of particular relevance

"E" earlier document but published on or after the international filing date

"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)

"O" document referring to an oral disclosure, use, exhibition or other means

"P" document published prior to the international filing date but later than the priority date claimed

"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention

"X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone

"Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art.

"&" document member of the same patent family

Date of the actual completion of the international search

25 July 2000

Date of mailing of the international search report

08/08/2000

Name and mailing address of the ISA

European Patent Office, P.B. 5818 Patentlaan 2
NL - 2280 HV Rijswijk
Tel. (+31-70) 340-2040, Tx. 31 651 epo nl,
Fax: (+31-70) 340-3016

Authorized officer

Meurisse, W

INTERNATIONAL SEARCH REPORT

International Application No
PCT/GB 00/01744

C.(Continuation) DOCUMENTS CONSIDERED TO BE RELEVANT

Category	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	KANG C -S ET AL: "A BROADBAND RING NETWORK: MULTICHANNEL OPTICAL SLOTTED RING" COMPUTER NETWORKS AND ISDN SYSTEMS,NL,NORTH HOLLAND PUBLISHING. AMSTERDAM, vol. 27, no. 9, 1 August 1995 (1995-08-01), pages 1387-1398, XP000519158 ISSN: 0169-7552	12
A	figure 3 ---	1-11
A	"Dynamic Packet Transport Technology and Applications Overview" CISCO SYSTEMS, 'Online! February 1999 (1999-02), pages 1-17, XP002143361 Retrieved from the Internet: <URL:http://www.cisco.com/warp/public/cc/t echno/media/wan/sonet/dpt/dpta_wp.htm> 'retrieved on 2000-07-25! cited in the application the whole document ---	1-12
A	WO 97 05725 A (TATHAM MARTIN CHRISTOPHER ;COTTER DAVID (GB); BRITISH TELECOMM (GB) 13 February 1997 (1997-02-13) page 26, line 13 - line 22; figure 23 -----	12

INTERNATIONAL SEARCH REPORT

Information on patent family members

International Application No

PCT/GB 00/01744

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
EP 0462349 A	27-12-1991	DE 69017203 D	30-03-1995
		DE 69017203 T	10-08-1995
		JP 2022732 C	26-02-1996
		JP 4233354 A	21-08-1992
		JP 7054939 B	07-06-1995
		US 5140587 A	18-08-1992
WO 9705725 A	13-02-1997	AU 712810 B	18-11-1999
		AU 6623496 A	26-02-1997
		CA 2228219 A	13-02-1997
		CN 1192307 A	02-09-1998
		EP 0872087 A	21-10-1998
		JP 11510335 T	07-09-1999
		NO 980366 A	20-03-1998
		NZ 313707 A	29-04-1999

ORIGINAL
NO MARSHALL

THIS PAGE BLANK (USPTO)